


CONNEXIONS



The Interoperability Report

May 1988

Volume 2, No. 5

*ConneXions -
The Interoperability Report
tracks current and emerging
standards and technologies
within the computer and
communications industry.*

In this issue:

SLIP: Serial Line IP.....	2
The Rapid Commercialization of TCP/IP.....	6
The Future of TCP/IP.....	10
The right ConneXions.....	15

From the Editor

In our first anniversary issue, we take another look at the TCP/IP marketplace. We do this in two articles. In the first, Michael Howard of Infonetics looks at the rapid commercialization of TCP/IP. The number of vendors supporting TCP/IP is growing every day, and there are no signs that this trend will change until OSI products are readily available.

In the second article, Bart Burstein (now with Evans & Sutherland) describes the role of TCP/IP coupled with NetBIOS in the Federal Government.

John Romkey is back again, this time with a description of SLIP: Serial Line IP. This protocol is receiving much attention because of its ability to provide low-cost internetworking over normal telephone circuits. With the advent of high-speed dial-up modems, this can provide a viable alternative to leased lines. I fondly remember my first acquaintance with SLIP in 1984 when we added the second Internet site in Norway by running SLIP (at 1200 baud!) between the University of Oslo and the Norwegian Telecommunications Administration Research Establishment. It was *slow*, but it worked!

Work is also underway to enhance SLIP. We will bring you more information about this as it emerges.

It is time once again to thank all those people who make *ConneXions* possible, first and foremost the contributors who provide valuable insight in their areas of expertise. As always, we are open to suggestions for future articles. A couple of Internet "gurus" have offered to help with a Questions & Answers column. So send us your ideas, questions, flames, user experiences, etc. Your input is always appreciated.

We have recently discovered (and rectified) a database problem which may have caused some irregularities in the delivery of *ConneXions* to a few of our subscribers. If you are missing back issues, please let us know and we will be happy to correct the situation.

ConneXions is published by Advanced Computing Environments, Inc., 480 San Antonio Road, Suite 100, Mountain View CA 94040, USA. Phone: 415-941-3399

© 1988

Advanced Computing Environments, Inc.
Quotation with attribution encouraged.

ISSN 0894-5926

SLIP: Serial Line IP

by John Romkey

Introduction

The TCP/IP protocol family runs over a variety of network media: IEEE 802.3 (Ethernet) and 802.5 (Token Ring) LAN's, X.25 lines, satellite links, and serial lines. There are standard encapsulations for IP packets defined for many of these networks, but there is no standard for serial lines. "SLIP" is currently a de facto standard, commonly used for point-to-point serial connections running TCP/IP. It is *not* an Internet standard.

History

SLIP has its origins in the 3Com UNET TCP/IP implementation from the early 1980's. It is merely a packet framing protocol: SLIP defines a sequence of characters that frame IP packets on a serial line, and nothing more. It provides no addressing, packet type identification, error detection/correction or compression mechanisms. Because the protocol does so little, though, it is usually very easy to implement.

Around 1984, Rick Adams implemented SLIP for UNIX 4.2BSD and Sun Microsystems workstations, and released it to the world. It quickly caught on as an easy, reliable way to connect TCP/IP hosts and routers with serial lines.

Application

SLIP is commonly used on dedicated serial links and sometimes for dialup purposes, and is usually used with line speeds between 1200bps and 19.2Kbps. It is useful for allowing mixes of hosts and routers to communicate with one another (host-host, host-router and router-router are all common SLIP network configurations).

Availability

SLIP is available for most Berkeley UNIX-based systems. It is included in the standard 4.3BSD release from Berkeley. SLIP is available for Ultrix, Sun UNIX and most other Berkeley-derived UNIX systems. Some terminal concentrators and IBM PC implementations also support it.

SLIP for Berkeley UNIX is available via anonymous FTP from `uunet.uu.net` in `pub/sl.shar.Z`. Be sure to transfer the file in binary mode and then run it through the UNIX `uncompress` program. Take the resulting file and use it as a shell script for the UNIX `/bin/sh` (for instance, `/bin/sh sl.shar`).

The protocol

The SLIP protocol defines two special characters: **END** and **ESC**. END is octal 300 and ESC is octal 333 (not to be confused with the ASCII ESCape character; for the purposes of this discussion, ESC will indicate the SLIP ESC character). To send a packet, a SLIP host simply starts sending the data in the packet. If a data byte is the same code as an END character, a two byte sequence of ESC and octal 334 is sent instead. If it the same as an ESC character, a two byte sequence of ESC and octal 335 is sent instead. When the last byte in the packet has been sent, an END character is transmitted.

Phil Karn suggests a simple change to the algorithm, which is to begin as well as end packets with an END character. This will flush any erroneous bytes which have been caused by line noise. In the normal case, the receiver will simply see two back-to-back END characters, which will generate a bad IP packet.

If the SLIP implementation does not throw away the zero-length IP packet, the IP implementation certainly will. If there was line noise, the data received due to it will be discarded without affecting the following packet.

Because there is no "standard" SLIP specification, there is no real defined maximum packet size for SLIP. It is probably best to accept the maximum packet size used by the Berkeley UNIX SLIP drivers: 1006 bytes including the IP and transport protocol headers (not including the framing characters). Therefore any new SLIP implementations should be prepared to accept 1006 byte datagrams and should not send more than 1006 bytes in a datagram.

Deficiencies

There are several features that many users would like SLIP to provide which it doesn't. In all fairness, SLIP is just a very simple protocol designed quite a long time ago when these problems were not really important issues. The following are commonly perceived shortcomings in the existing SLIP protocol:

- *Addressing*: both computers in a SLIP link need to know each other's IP addresses for routing purposes. Also, when using SLIP for hosts to dialup a router, the addressing scheme may be quite dynamic and the router may need to inform the dialing host of the host's IP address. SLIP currently provides no mechanism for hosts to communicate addressing information over a SLIP connection.
- *Type identification*: SLIP has no type field. Thus only one protocol can be run over a SLIP connection, so in a configuration of two DEC computers running both TCP/IP and DECnet, there is no hope of having TCP/IP and DECnet share one serial line between them while using SLIP. While SLIP is "Serial Line IP", if a serial line connects two multi-protocol computers, those computers should be able to use more than one protocol over the line.
- *Error detection/correction*: noisy phone lines will corrupt packets in transit. Because the line speed is probably quite low (likely 2400 baud or less), retransmitting a packet is *very* expensive. Error detection is not absolutely necessary at the SLIP level because any IP application should detect damaged packets (IP header and UDP and TCP checksums should suffice), although some common applications like NFS usually ignore the checksum and depend on the network media to detect damaged packets. Because it takes so long to retransmit a packet which was corrupted by line noise, it would be efficient if SLIP could provide some sort of simple error correction mechanism of its own.
- *Compression*: because dialin lines are so slow, packet compression would cause large improvements in packet throughput. Usually streams of packets in a single TCP connection have few changed fields in the IP and TCP headers, so a simple compression algorithm might just send the changed parts of the headers instead of the complete headers.

Some work is being done by various groups to design and implement a successor to SLIP which will address some or all of these problems.

SLIP: Serial Line IP *(continued)*

SLIP drivers

The following C language functions send and receive SLIP packets. They depend on two functions, `send_char()` and `recv_char()`, which send and receive a single character over the serial line.

```

/* SLIP special character codes
 */
#define      END      0300    /* indicates end of packet */
#define      ESC      0333    /* indicates byte stuffing */
#define      ESC_END  0334    /* ESC ESC_END means END data byte */
#define      ESC_ESC  0335    /* ESC ESC_ESC means ESC data byte */

/* SEND_PACKET: sends a packet of length "len",
 * starting at location "p".
 */
void send_packet(p, len)
    char *p;
    int len; {

    /* Send an initial END character to flush out any data that
     * may have accumulated in the receiver due to line noise
     */
    send_char(END);

    /* For each byte in the packet, send the appropriate
     * character sequence
     */
    while(len--) {
        switch(*p) {
            /* If it's the same code as an END character, we
             * send a special two character code so as not
             * to make the receiver think we sent an END
             */
            case END:
                send_char(ESC);
                send_char(ESC_END);
                break;

            /* If it's the same code as an ESC character, we
             * send a special two character code so as not
             * to make the receiver think we sent an ESC
             */
            case ESC:
                send_char(ESC);
                send_char(ESC_ESC);
                break;

            /* otherwise, we just send the character
             */
            default:
                send_char(*p);
        }
        p++;
    }

    /* Tell the receiver that we're done sending the packet
     */
    send_char(END);
}

/* RECV_PACKET: receives a packet into the buffer
 * located at "p". If more than len bytes are
 * received, the packet will be truncated.
 * Returns the number of bytes stored in the buffer.
 */
int recv_packet(p, len)
    char *p;
    int len; {
    char c;
    int received = 0;

```



```

/* Sit in a loop reading bytes until we put together
 * a whole packet.
 * Make sure not to copy them into the packet if we
 * run out of room.
 */
while(1) {
    /* Get a character to process
     */
    c = recv_char();

    /* Handle bytestuffing if necessary
     */
    switch(c) {

        /* If it's an END character then we're done
         * with the packet
         */
        case END:
            /* A minor optimization: if there is no data
             * in the packet, ignore it. This is meant
             * to avoid bothering IP with all the empty
             * packets generated by the duplicate END
             * characters which are in turn sent to try
             * to detect line noise.
             */
            if(received)
                return received;
            else
                break;

        /* If it's the same code as an ESC character, wait
         * and get another character and then figure out
         * what to store in the packet based on that.
         */
        case ESC:
            c = recv_char();

            /* If "c" is not one of these two then we
             * have a protocol violation. The best bet
             * seems to be to leave the byte alone and
             * just stuff it into the packet
             */
            switch(c) {
                case ESC_END:
                    c = END;
                    break;
                case ESC_ESC:
                    c = ESC;
                    break;
            }

        /* Here we fall into the default handler and let
         * it store the character for us
         */
        default:
            if(received < len)
                p[received++] = c;
            }
    }
}

```

JOHN ROMKEY received his B.S. in Computer Science from MIT in 1985. While there he worked on the PC/IP project and later left to help found FTP Software. He is currently applying for membership to the Bio-of-the-Month Club!

The Rapid Commercialization of TCP/IP

by Michael Howard, Infonetics, Inc.

The rapid commercialization of TCP/IP is a market phenomenon with its roots in the accelerating demand for internetworking. What is internetworking? What are the market pressures driving the demand for internetworking? What do we mean by "TCP/IP", and how can we measure its growth? Is TCP/IP the only answer? And, finally, what does the future hold? In this article I will attempt to answer these questions and give you Infonetics' view of this important developing commercial marketplace.

What is internetworking?

Internetworking is the ability for users to communicate across different networks. Actually, it is the ability for computers of various types to communicate across different networks. Networks of different types can be connected by gateways, such as a DECnet-SNA gateway. The gateway is a method of translating between dissimilar networks but is not an efficient method of internetworking. Using gateways is similar to working in an office where one workgroup speaks French and another speaks English. In order for the two workgroups to communicate, or share information, an interpreter is required for every transaction.

Multi-vendor networks

The internetworking that we are discussing here is the internetworking needed in the marketplace. An internetworking where gateways are not needed, but rather the networks are connected by bridges or routers. It is also multivendor, another reality in the marketplace. If an organization is dedicated to a single vendor such as DEC or IBM, then there is no problem for internetworking: you just buy the vendor's solution when it is available and you buy it at the vendor's price.

The main forms of internetworking communication are electronic mail, file transfer, and terminal emulation (remote login). These applications resulted from the demand for internetworking, and are the way that users can access, manipulate, move, and exchange information. Not surprising then, the demand for internetworking is based on the need and desire for access to information that resides on other PCs, departmental processors, and mainframes. When users have computer workstations on their desks and know that information they need resides on a computer elsewhere, then the desire for access to that information creates demand for internetworking.

What is driving the demand for internetworking?

The events leading to today's demand for internetworking began with the explosion of personal computers in the workplace in the early 1980s. The PC became commonplace on the desktop as a personal productivity tool. PC users found that it was more productive to share spreadsheet data, work on reports or memos together, and share other information, thus creating the concept of the workgroup. This type of information exchange was executed by carrying a floppy diskette from one PC to another ("Sneaker Net"). PCs were also outfitted with terminal emulation cards and software to access the departmental mini or the corporate mainframe.

Local Area Networks (LANs) were introduced and allowed the next stage of productivity to occur at the workgroup level. Sneaker net was replaced within workgroups by a PC LAN. Today, LANs are the fastest growing segment of the computer industry. That's the good news. The bad news is that there are many kinds of LANs that don't speak the same protocols, and so the desire to exchange information between workgroups is unsatisfied.


Most workgroup networking environments are multivendor. The most popular PC LANs are by IBM, Novell, 3Com and Apple. These PC LANs use gateways to connect to minis, mainframes, and other PC LANs. Other LANs were designed to connect terminals to various computers, but not particularly PCs. These LANs include DECnet and UNIX networks, as well as those produced by Ungermann-Bass, Excelan, and Bridge. The multivendor environment does not lead to easy solutions for internetworking.

Major internetworking solutions

So, when we consider multivendor internetworking, the choice is reduced to two solutions: (1) the promised ISO/OSI, and (2) TCP/IP. The OSI products are just becoming available. (See "Will OSI Ever Become a Reality?" in the February 1988 issue of *ConneXions*). TCP/IP has been strong in the defense, university, and UNIX market segments for many years. This strength is due mainly to two factors: TCP/IP is a DoD requirement, and is the basic networking protocol of UNIX. The development of TCP/IP with the Arpanet in the 1970s was based on the requirements of providing access to multivendor computing environments. The choice of TCP/IP for UNIX networking was based largely on its availability as public domain software funded by our tax dollars.

Measures of TCP/IP growth

There are several measures of the rapid growth of TCP/IP. One easy metric is the number of vendors offering TCP/IP products as shown in the chart below. The numbers are rough, but fairly accurate, and are drawn from knowledgeable TCP/IP industry participants. Notice the classic "hockey stick" exponential growth curve.



continued on next page

Rapid Commercialization of TCP/IP (continued)


Another measure is the existence of Advanced Computing Environments. This company is providing a focal point for the TCP/IP marketplace and for practical improvements to the TCP/IP product set, which are necessary for growth in the commercial sector. A focal point means that a forum is provided for users and vendors to agree on what is needed, so that vendors can produce products that "interoperate". Interoperation means that products from different vendors on the same network will work together -- a requirement for commercial growth of the TCP/IP market.

NetBIOS

A good example of the results of these forums is the implementation of NetBIOS on top of TCP/IP (NetBIOS is the de facto standard networking software from IBM). Several vendors demonstrated the interoperation of NetBIOS capability on TCP/IP at the 2nd TCP/IP Interoperability Conference in December of 1987. Another area being worked on is network management, a much tougher problem than NetBIOS.

Conference attendance

The increase in commercial interest in TCP/IP can be measured by looking at the attendance at the 1st and 2nd TCP/IP Interoperability Conferences put on by Advanced Computing Environments in April and December of 1987. The next chart shows how the attendance by the commercial sector has increased remarkably, while the vendor and research communities have remained relatively constant.



There is a decided lack of market research data on the TCP/IP marketplace. None of the major computer industry market research organizations have published detailed information. Infonetics has developed a TCP/IP market forecast for the commercial sector through 1990.

Today the TCP/IP market is roughly 65% defense and university, and 35% in the commercial sector. By 1990, this fraction will reverse, with 65% of the revenues in the commercial sector. Infonetics projects that the worldwide commercial market will be \$540M in 1990. Infonetics and Advanced Computing Environments have undertaken a major effort in this area to produce two market research products that will fill this need for market data.

The future

So, what does the future hold for TCP/IP? Infonetics obviously believes that the TCP/IP market is on a rapid increase, but what does that mean to the commercial marketplace? We can expect some of the usual changes that come with an expanding marketplace. There will be more products from more vendors. The products will become faster and will cost less. There will be more computing systems with TCP/IP availability.

UNIX

One factor in the TCP/IP market is the fact that UNIX systems use the TCP/IP protocols for networking. Since UNIX is a strong part of the TCP/IP marketplace, we see that the movement toward a single standard UNIX in the next two years will also be a growth factor. Standardization usually leads to market growth.

As OSI gradually comes to the marketplace, it remains to be seen what will happen to TCP/IP. Organizations with TCP/IP products in place are not likely to buy replacement equipment without some clear benefit. That clear benefit is not apparent today. There are efforts underway to provide the common OSI applications, such as X.400, FTAM, and CASE, on top of TCP/IP so that OSI applications will be available on TCP/IP networks.

Multiple stacks

The trend that Infonetics sees is that vendors will provide multiple protocol stacks on the same networks, for example, TCP/IP and the OSI transport protocol TP-4 (and the OSI IP) running on Ethernet. Some vendors have PC network cards with memory for the transport protocols in software. As OSI becomes available, the plan is to offer new software without the necessity of new hardware. Excelan has announced a product that allows Novell Netware's IPX protocol to co-exist on their TCP/IP network, allowing access to Novell LAN file servers. DEC is talking about DECnet Phase 5, which would allow choice of OSI, TCP/IP, or DECnet protocols over the same Ethernet. IBM will continue to treat SNA as its networking product, but allow gateways to other types of networks. Finally, Apple Computer is working with Ungermann-Bass to produce TCP/IP for the Macintosh.

So, we see that TCP/IP is growing fast, especially in the commercial sector. Several factors are helping its growth, mainly a demand for internetworking and the lack of other suitable solutions. Infonetics has some measures of the recent growth, and will have new market research numbers available in May 1988. Over the next three to five years, we see continued growth. The unknown beyond 1990 is whether the OSI window of opportunity will close -- only time will tell.

MICHAEL HOWARD is executive vice president and senior analyst for internetworking at Infonetics, Inc. Infonetics is an international consulting company providing clients with publications, custom research and analysis services, and industry conferences, with offices located in Santa Clara, California.

The Future of TCP/IP

by Bart Burstein

The reports of TCP/IP's death are premature. Although fifteen government agencies disclosed in 1986 that they would band together to develop and implement networking based on the International Organization for Standardization's Open Systems Interconnection (ISO/OSI) reference model, practical availability of those products is still many years away. Many users in the federal government and elsewhere desire a move to ISO/OSI protocols right away. From a technological point of view, however, such a move is not only infeasible but also potentially unwise. The ISO/OSI suite of multi-vendor networking protocols still needs time to mature. Consequently, TCP/IP will endure as a federal government networking standard for many years to come.

Integration of PCs

Of more immediate concern to the federal government networking community is the necessity to integrate the exploding population of personal computers into existing federal government networks. The growing importance of PCs in the government is leading to what amounts to a marriage between the NetBIOS networking interface, the standard in PC networking, and the TCP/IP protocol suite.

How it started

TCP/IP is the most widely known set of protocols of the Department of Defense's (DoD) Internetworking Protocol Suite, also known simply as the Internet Protocols. The DoD developed the Internet Protocol suite to meet its own special needs. The Internet Protocol Suite, used by the Defense Data Network (DDN) and other federal government networks, consists of many protocols, including the Internet Protocol (IP), Transmission Control Protocol (TCP), Telnet, a virtual terminal protocol, File Transfer Protocol (FTP), and Simple Mail Transfer Protocol (SMTP). Of these, TCP and IP have become closely identified with each other and are widely used together, both inside and outside the government, especially in the scientific, academic, engineering, and research communities that do a good deal of basic research for the DoD.

NetBIOS

NetBIOS, in contrast, originated in the commercial world. NetBIOS, which stands for the Network Basic Input/Output System, was created by IBM for the express and sole purpose of permitting stand-alone PCs to communicate with each other and to share data, resources, such as printers, and application programs over a local area network. A simple enough mission, and one for which NetBIOS is well suited.

Because of IBM's backing and its ability to set the course of the IBM PC standard architecture's evolution, NetBIOS has become the de facto networking interface standard for today's office. It is also the standard to which today's most popular, sophisticated, and useful office applications software is being written.

Why NetBIOS over TCP/IP?

Federal government agencies and departments have been experiencing the same rapid influx of desktop PCs that their commercial counterparts have. In response, many federal agencies and departments have embarked on large-scale procurements that will place advanced PCs on practically every desktop in many government settings.

Large procurements

This move toward the wide-scale use of PCs is especially apparent in the various branches of the DoD. Large contracts for tens of thousands of PCs have been awarded to PC suppliers in the last two years. The U.S. Air Force, for instance, is in the midst of a long-term procurement from Zenith Data Systems Corp. that will eventually comprise between 80,000 and 100,000 IBM PC AT-compatible workstations. This procurement has proven so successful and the workstations so popular, in fact, that other branches are following suit.

Applications

These PCs will be used to perform a variety of desktop workstation applications, such as word processing and spreadsheets. That means they will need to have access to the large base of popular PC-based applications programs, such as Lotus 1-2-3, Multimate, Microsoft Word, and others. Virtually all of these programs support multiuser access in LAN environments through the NetBIOS interface.

Additionally, these PCs will often replace existing terminals that are now linked to minicomputers or mainframes over DoD networks; sometimes PCs will be installed rather than adding new terminals to existing systems and networks.

Resource sharing

One of the primary requirements of these PCs will be that they be able to share resources and peripherals not only with each other, but also with larger minicomputer and mainframe systems on existing networks. They also must have access to existing federal government electronic mail networks. Support of the TCP/IP protocol suite, therefore, is essential.

Finally, NetBIOS over TCP/IP matches the general movement in the DoD and other federal agencies to support "standard" components for computer and networking procurements. These commercial off-the-shelf (COTS) components include the AT-compatible desktop workstation from Zenith (dubbed the Z-248). Unix-based, multiuser supermicros, and, now, networking in the form of NetBIOS over TCP/IP that can use existing MILNET access.

Benefits of NetBIOS over TCP/IP

A session-level interface standard defined by IBM and Microsoft, NetBIOS describes the services necessary for PC communications in a LAN environment. Implementing NetBIOS over TCP/IP has two primary benefits. First, NetBIOS over TCP/IP enables PCs to function as equal members in multivendor computing networks. Second, it standardizes the underlying communications protocols, which in turn creates a standard programming environment in the upper layers of the protocol stack.

Having a standard interface unlocks applications written to that interface from the underlying layers of network interfaces, services and protocols. This, in essence, will enable users to move applications written to the NetBIOS interface standard from one underlying protocol set to the next, as need develops.

The newfound independence will set the stage for eventual transition to OSI. As long as the interface remains constant, users can move from one set of services and protocols to another without disrupting their software investments.

continued on next page

The Future of TCP/IP (*continued*)

The NetBIOS over TCP/IP standard also opens the existing TCP/IP network environment to the industry standard for PCs, which in turn opens up PC networking and workstations to most government networks.

An office standard

A third reason for the popularity of the NetBIOS over TCP/IP standard is that NetBIOS is the networking interface standard for the office. Many office applications have already been written to that standard. Running NetBIOS over TCP/IP gives the existing TCP/IP user community access to the most popular PC applications available today, such as Lotus 1-2-3, Microsoft Word, MultiMate, and others written to that standard today, as well as those that will be written to it in the future.

Finally, NetBIOS is an "off-the-shelf" standard. Users can buy NetBIOS applications at a local retail store. This is in keeping with the movement in federal government computer procurements to COTS products and standards.

Technical issues

Some key issues had to be resolved by the standards group, as many technical decisions were required in marrying NetBIOS services to the TCP/IP environment. For instance, TCP is a byte-oriented protocol, while NetBIOS sessions provide a message-oriented interface. A standard way had to be worked out to match one to the other.

Naming is also an issue, because NetBIOS was conceived as a network interface for small, departmental LANs, many of which consists of a dozen users or less. TCP/IP, on the other hand, was conceived to work on networks with literally thousands and thousands of users. It, therefore, has network *addresses*. Because NetBIOS has no concept of what a network address is, a mechanism is needed to convert names to addresses.

Other issues include internetworking, or communicating between different physical networks, and multiple recipient support (multicasting). This latter capability is still a research topic in the TCP/IP world and does not have a formalized, approved specification.

Beyond NetBIOS - TCP/IP and OSI

Despite a growing movement toward NetBIOS over TCP/IP, government agencies, particularly the DoD, are still pledged to pursuing the ISO/OSI suite of multi-vendor networking protocols. However, the various components of that protocol suite are still under development. In the opinion of many observers, the ISO/OSI suite is at least five years away from widespread acceptance and practical implementation in the government computer user community. Then a transition period must follow. It may be ten to fifteen years before TCP/IP dies away. In the meantime, TCP/IP and NetBIOS will fill the gap.

Valuable lessons

But more than that, since NetBIOS over TCP/IP frees the applications from the underlying interfaces, services, and protocols, it will prepare the way for ISO/OSI, in terms organizational experience and procedures. Moreover, the current transition now under way to NetBIOS over TCP/IP within the government computer community will provide many valuable lessons for the eventual transition to OSI.

Users will need experience in transitioning if the eventual move from TCP/IP to ISO/OSI is to prove successful. NetBIOS over TCP/IP serves as a good way to get that experience.

The history of a proposed standard

The linking of NetBIOS to TCP/IP has become so important for the future of government computer usage that a group of LAN and communications vendors have banded together for the dual purposes of defining a standard specification and accelerating the development of NetBIOS over TCP/IP usage. In this way, it is hoped that this standards group can prevent the proprietary, and therefore, incompatible implementations of NetBIOS over TCP/IP.

ULANA's contribution

The NetBIOS over TCP/IP standardization effort stems from work done by the Mitre Corporation, a federal government contractor, for the U.S. Air Force's Universal Local-Area Network Architecture (ULANA) project. ULANA is one project of a three-part effort by the Air Force to acquire common "building block" components for computer usage and procurements with the service, the other two being AT-compatible desktop workstations and UNIX-based multiuser super-micros. Together, these three components will make up a common hardware and networking platform for application development within the Air Force computer community.

The Air Force chartered Mitre to develop ULANA's specifications, which the service planned to use as the basis of future Air Force networks. Originally, ULANA was conceived as a proprietary, Air Force Specified LAN. At that time, several years ago, no commercially available networking products could meet the combined needs for secure, high-speed links that could connect both asynchronous and synchronous devices to computers with more than 20 different host operating systems.

In 1986, however, Mitre and the Air Force realized that many of the features they desired in ULANA would soon be readily available in the marketplace from commercial suppliers. It was then that Mitre and the Air Force abandoned work on the proprietary interface in favor of a commercially available, "standard" approach.

For that standard, Mitre decided that the ULANA specification should combine the TCP/IP protocol suite, universally accepted within the Arpanet and MILNET communities, with the NetBIOS interface. NetBIOS won out over other offerings available in the commercial marketplace, including Sun Microsystems' Network File System (NFS) and AT&T's UNIX System V-based Remote File System (RFS). NetBIOS was selected because it has more widespread support in the commercial marketplace, due in part to the availability of applications software.

Vendors meet

Mitre and its supporters, quickly realized that the number of NetBIOS over TCP implementations were rapidly proliferating, a development that would create serious incompatibility problems for federal government users. Under the auspices of Mitre, a group was formed to develop a specification for running NetBIOS over TCP/IP. This specification would be proposed to the vendor and user community as an industry standard. Vendors would then use this specification as a base for their own implementations, which may or may not include proprietary enhancements and additions.

continued on next page

The Future of TCP/IP (continued)

Two RFCs

In March of 1987, after almost a year of development, Mitre held a meeting in Monterey, California, for interested vendors. Those attending included LAN vendors, and representatives from several workstation and minicomputer suppliers. That meeting resulted in the publication of two Request For Comments (RFCs): RFC 1001 and RFC 1002. Publication of an RFC is the recognized procedure for establishing new communications standards in the Arpanet and MILNET [or "Internet"] communities. However, it is not the end of the process.

Currently, the two RFCs for the NetBIOS over TCP/IP standard are being circulated in draft proposal form for comment. Vendors may use the draft proposal RFCs as the basis for their respective real-world products that implement the NetBIOS over TCP/IP standard.

[Reprinted with permission from *Government Data Systems*, September, 1987. Copyright by Media Horizons, Inc., 1987]

BARTON BURSTEIN is currently Product Marketing Manager with Evans & Sutherland Computer Division in Mountain View, California. This article was written while he was Net/One TCP Product Line Manager with Ungermann-Bass, where he championed the design, implementation and marketing of a TCP/IP oriented product line. Burstein has a broad range of marketing experience in communications and networking products. He has also been a successful systems programmer, systems architect and manager. He has been an officer of an ANSI committee involved with OSI standards, and holds masters and bachelors degrees from the University of Michigan and Antioch College.

Here we are!

Like the OSI model, Advanced Computing Environments now counts seven members. Shown here in front of our new headquarters in Mountain View are (left to right): *Dan Lynch*, President, *Margot Lockwood*, Director of Conferences, *Darlene Billstrom*, Business Manager, *Ole Jacobsen*, Editor, *Wendy Gibson*, Director of Sales and Marketing, *Sara Tietz*, Director of Education, *Valarie Collins*, Assistant Business Manager.



Getting the right ConneXions


Choosing a name for a product or service is always difficult, and although we went through all the correct legal steps and a great deal of creative thinking, inevitably there will be other items with a similar name. In my travels, I have come across some of these, and below is my current collection. Remember, this newsletter is *ConneXions - The Interoperability Report*, ask for the original! --Ole




AUTHORISED MOTOROLA DEALER

No more missed calls – keep your own phone and private number by your side, wherever you are. Full international dialling and receiving facilities available, even in your car or on the train!

British cellular telephone rental company



British Rail Travel Service




Michigan Bell newsletter

CONNECTIONS

The Newsletter for Networking Macintoshes

Macintosh Networking newsletter



Halley Systems product news



The Interoperability Report

...and this is it!

CONNEXIONS

480 San Antonio Road
Suite 100
Mountain View, CA 94040

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

CONNEXIONS

PUBLISHER Daniel C. Lynch

EDITOR Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf, Vice President, National Research Initiatives.

Dr. David D. Clark, The Internet Architect, Massachusetts Institute of Technology.

Dr. David L. Mills, NSFnet Technical Advisor; Professor, University of Delaware.

Dr. Jonathan B. Postel, Assistant Internet Architect, Internet Activities Board; Associate Director, University of Southern California Information Sciences Institute.

Subscribe to CONNEXIONS

U.S./Canada \$100. for 12 issues/year
International \$ 50. additional per year

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

☐ Check enclosed (in U.S. dollars made payable to CONNEXIONS). ☐ Bill me/PO# _____

☐ Charge my ☐ Visa ☐ Master Card Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to:
Back issues available upon request \$10./each

CONNEXIONS

480 San Antonio Road Suite 100
Mountain View, CA 94040
415-941-3399

CONNEXIONS